U.S. NAVAL ACADEMY
COMPUTER SCIENCE DEPARTMENT
TECHNICAL REPORT

Resource Contrained Network Design and Implementation

Gawne, Michael B.

USNA-CS-TR-2009-02

January 3, 2009

USNA Computer Science Dept. ◦ 572M Holloway Rd Stop 9F ◦ Annapolis, MD 21403

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE
03 JAN 2009

2. REPORT TYPE

3. DATES COVERED
00-00-2009 to 00-00-2009

4. TITLE AND SUBTITLE
Resource Contrained Network Design and Implementation

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
U.S. Naval Academy,Computer Science Department,572M Holloway Rd Stop 9F,Annapolis,MD,21403

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR’S ACRONYM(S)

11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**25**	

U.S. NAVAL ACADEMY
COMPUTER SCIENCE DEPARTMENT
TECHNICAL REPORT

Resource Constrained Network Design and Implementation

Gawne, Michael B

USNA-IT

8 December 2008

Computer Science Department
SI495B: Research Project Report
Spring AY08

Resource Constrained Network Design and Implementation

by

Midshipman Michael Gawne

United States Naval Academy
Annapolis, MD

____

____
Date

Certification of Faculty Mentor's Approval

Thomas Augustine, Lt Col, USAF
Department of Computer Science

____

____
Date

Department Chair Endorsement

Thomas Logue, CAPT, USN
Chair, Department of Computer Science

____

____
Date

Table of Contents

Table of Figures

Table of Tables

Executive Summary

The research study, *Resource Constrained Network Design and Implementation*, was conducted at the United States Naval Academy (USNA) in order to simulate creating a secure and functional network with constrained resources. This is a problem faced by any administrator trying to maximize the security and functionality of a private network while minimizing the costs involved in the creation and maintenance of a local area network (LAN).

Computers are utilized for everything from electronic correspondence via e-mail, to multimedia entertainment, to developing simulations and models of complex engineering designs. However computers are far less useful when isolated off of a network. In this modern age, the ability to network computers and other electronic devices is nearly a requirement in all industries. While the ability to network is essential, the network must be designed in a manner providing confidence in the security of the system and the data.

The goals of this research study were established in a step-wise foundation, each step relying upon the completion of the previous goal. The goals were: 1) Configure the most basic LAN, one comprised of a switch and two nodes and establish communication between the two nodes, 2) Introduce a 3600 series router into the network dividing the two LANs to simulate an internal network LAN, and the external internet and permit communications between the two networks, 3) Create Virtual LANs (VLAN) and implement their use on the network, 4) Utilize Access Control Lists with rulesets based upon the VLANs, 5) Introduce a PIX 515E firewall into the network to isolate and protect the internal LAN from the simulated internet.

At the conclusion of the study it was deemed that education and knowledge levels must be included when considering all the resources available to the network administrator. By far, the most difficult segment of this research project was the last segment, introducing the PIX 515E firewall into the network and establishing communications through the firewall. Internet resources are plentiful in the form of forums, web blogs, and other information technology support sites for novice network administrators to educate themselves on the basics of configuring their network, but online education for establishing a firewall is not as plentiful. Furthermore, even with the assistance of two Cisco Certified Network Associates (CCNA), a mid-level certification, and a Cisco Certified Internetwork Expert (CCIE), it is sometimes difficult to find the one error that drops all traffic and renders your network isolated from the rest of the world. One of the most beneficial ways to spend precious resources may very well be a seminar or other education venue where the network administrator may first hone his/her skills on firewall configuration before attempting to configure the firewall with no formalized training. Through the review of this research study it appears that while a secure, effective network can be constructed in a resource constrained environment, especially when resources are properly allocated.

Ultimately I accomplished the goals of this study through literature review and experimental setups. Throughout this study we conclude that while many organizations design networks based on a financial budget, the greatest constraint is the knowledge of those individuals who will maintain the network on a daily basis. I recommend that system administrators and others who maintain the network should exercise significant input in the network design process. Their experience and knowledge should both be considered, as their experience may optimize the network, and a lack of knowledge may require additional education before enacting a network design. The money spent on their education typically is an investment providing yields far greater than the initial investment, and their experience may also prevent poor choices which will cost more time and money to remedy later.

Goals	Methods Used
Permit communications between two computers	Configure the most basic LAN, (2 nodes connected by a switch)
Connect two separate LANs	Configure a Cisco 3600 Series router to route traffic between the two LANs
Isolate network into compartments to filter extraneous traffic	Establish Virtual LANs.
Only permit standard traffic between VLANs	Configure Access Control Lists to prohibit unexpected traffic traveling between VLANs.
Enable Network Address Translation, protocol inspection, and traffic filtering to shield the internal network from the simulated internet	Introduce a PIX firewall into the network between the 3600 router and the internal LAN, making the connection between the 3600 and the firewall the access-point to the external internet.

Table 1 – Initial Goals

Laying the Foundation

When designing a network in a resource constrained environment, it is important for one to consider everything available to them before beginning the task. Even with unlimited resources it would never be possible to have a completely secure network, but squandering away any extra resources only increases the difficulty of securing a network. This makes it important to first consider the requirements of the network, and then to direct the majority of the resources into the largest risk area.

This research project was designed to simulate the typical small to medium-sized network design. That is to say a network with one point of access to the rest of the world wide web, a private network of workstations behind the access-point which requires the ability to access external internet sites, and servers running services which must accept incoming connections from both the internal intranet, and the external internet. This sort of network does not require the complexity of multiple exterior gateway routers, yet provides the necessary functionality to provide services to multiple users on either side of the access point.

Fiver Layer Network Model

This paper assumes the reader is familiar with the basic concepts of IP addressing, network addresses, subnet masks, and host addresses. The user should also understand the TCP/IP model of networking, as it is applicable when discussing the firewalls later. (Table 2 – Five Layer Model)

Application	The actual program or service (SMTP, SSH, POP3, IMAP, MySQL, HTTP, DNS, etc)
Transport	Transmission Controlled Protocol (TCP) or User Datagram Protocol (UDP)
Internet	Internet Protocol, ICMP
Link	Ethernet/FastEthernet
Physical	Electrical or Optical pulses (1's & 0's) in the medium

Table 2 – Five Layer Model

The services, or applications, that run on the typical network typically include DNS (Domain Name Server), POP3 or IMAP (Post Office Protocol 3 or Internet Message Access Protocol - Email Protocols), HTTP (Hyper Text Transfer Protocol – Web Server), and some form of a database service, MySQL in this example. Many Windows networks will also contain an ActiveDirectory server which allows for user rights and privileges to be stored centrally and affect the entire network.[1]

Network Layout and Design

After understanding the 5-layer networking model, the network administrator must determine what services will be necessary for the network to properly function. An e-commerce site will need some sort of database and web server to be included in its design; however an internet cafe will not likely need these, but may need an Active Directory server to maintain strict user privileges to maintain secure, untampered computers for the public to use. It is important to analyze each intended use of the network as a whole and make a list of all expectations required of the network. By analyzing the needs of the organization employing the network and properly designing the network in the beginning, the

[1] Odom, Wendel 2008. CCENT/CCNA ICND1: Official Certification Guide. 696 pages. ISBN:978-1-58720-182-0 Pages 16-21

network administrator effectively manages his resources by not placing hours or days of work into one direction only to realize that his design was flawed and he must restart the entire process of configuring the network.

In the example for this project, (Figure 1- Network Diagram) MySQL, Apache (HTTP), and PHP were the services required of the network. The network was also expected to provide workstations where users could access content from the internal intranet or the external internet. At this point the designer must determine how to logically lay out the network, that is how to configure the addressing scheme. An experienced network designer will compartmentalize the network into smaller segments, as it is easier to restrict the flow of unnecessary data between segments than between individual nodes on the network. Care must be taken to allow for future expansion of the network . In the case of a library computer lab, it would be foolish to design a subnet only allowing for 14 ip addresses in an area with 12 computers, because it is likely that in time the needs of the organization may change, thus requiring the network administrator to place more than 14 nodes in the computer lab. Placing more than 14 nodes in a configuration like this would require the administrator to completely reassign a network address for the computer lab as well as host addresses for each of the individual nodes in the lab. While it would be foolish to leave addressing space for over 65,000 nodes in the room, it would likely be wise to leave addressing space for 22 or 30 nodes in that computer lab at a minimum.[2]

[2] Ibid; 337-47

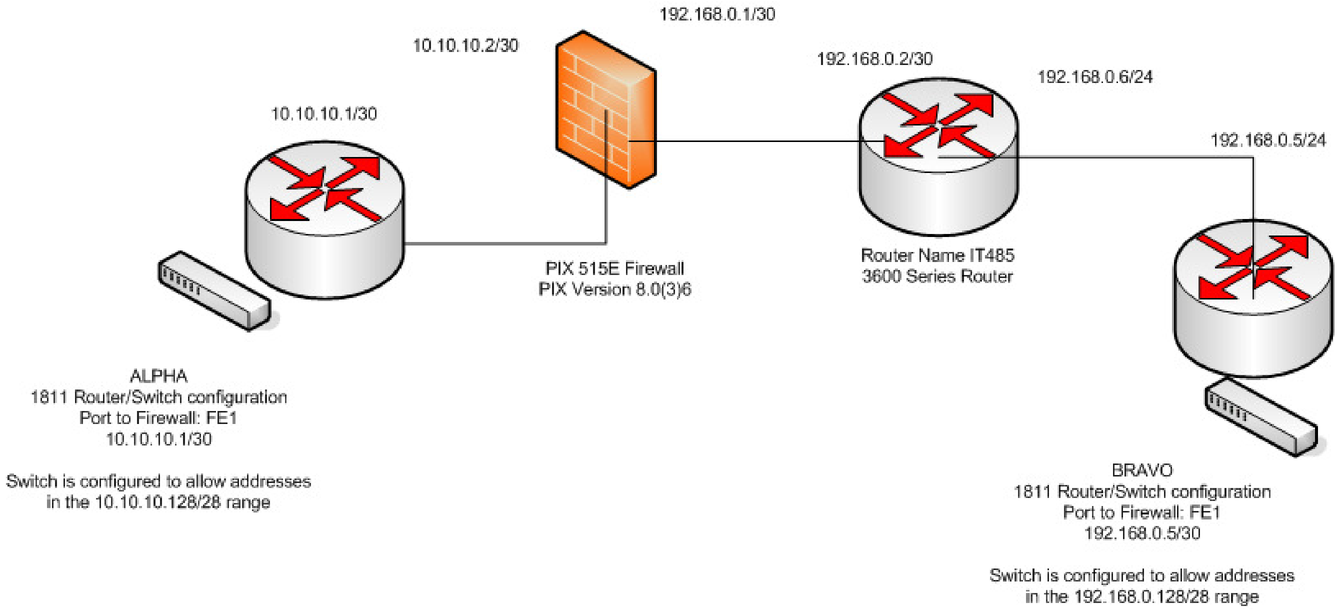

Figure 1- Network Diagram

Network Configuration

Up until this point in the design of the network, the work has been theoretical. That is to say no cables have been connected, no processors turned on, an no routers have been configured. The steps may have seemed to progress quickly and easily due to the simplicity of the network being demonstrated. Much like any parent tries to instill upon their children that prior planning prevents poor performance, network design is no exception. By properly taking these few steps to methodically analyze the requirements, a designer can reduce their risk of dedicating time and resources towards one method of pursuing a solution to only find that their work has been in vain.

Merely configuring a network in theory will not allow for the data to ever be passed across the medium. At some point the equipment needs to be turned on, cables connected, and hardware configured. In the case of the sample network, a Cisco 1811 was used, however any switch that can be configured to use VLANs, (for example the Cisco Catalyst 2950SX-48 SI), would have sufficed. If the network administrator needs to purchase new equipment, it is important that they find a switch which provides enough ports for every node on the network, plus sufficient unused ports for expansion. For example, the Catalyst 2950SX-48 SI is a network switch with 48 ports that typically retails for around $3000 new. It would likely be unwise to purchase the 2950SX-48 SI for a network that is only comprised of seven nodes, with little or no expected expansion. The administrator must analyze the needs of the network and determine which specific models of equipment to utilize.[3]

Configuration Techniques

Once the network design has been finalized, the network can begin to take form. When implementing a new network, the designer may take one of two methods.

In the first method, the designer may begin work on configuring the network starting from the single access point to the internet, (typically a router, but on small networks this may be a switch or a firewall), and initially configure a path to a single node. After ensuring that the single node has a path to the external internet, the administrator will then configure the rest of the nodes on that subnet, then continue the process for each of the paths to every node.

In the second method, the network administrator begins the work on a single subnet, ensures connectivity between all nodes on the subnet, and then moves up the path towards the access point to the internet. After configuring all equipment in the path up to the internet accesspoint, the administrator begins from the next leaf node of the tree.

Initial Configuration Method

The results of this study found that the most effective way to configure a brand new network was to combine the two methods mentioned above. In attempts to isolate configuration errors, its best to be able to connectivity from one leaf node to another leaf node in the same subnet. After ensuring that intra-subnet connectivity exists among two separate subnets, connectivity was established by

[3] CDW Product Overview: Cisco Catalyst 2950-SX-48SI 48 port Standalone Layer 2 Managed Switch. 2008. http://www.cdw.com/shop/products/default.aspx?EDC=553700

connecting the two switches to a router and configuring the router to forward the appropriate traffic between the two subnets.

In using this method, it is possible to ensure first that the computers (leaf nodes) of the network are able to communicate with each other, and verify that firewalls are disabled so that ICMP pings can be used to test connectivity between the nodes. Assigning static IP addresses to the leaf nodes in accordance with the addressing scheme is important for the initial phases of the configuration, as this can later be changed to utilize a Dynamic Host Configuration Protocol (DHCP) server if so desired. (Figure 2- Manually Assigning IP Addresses in Windows)

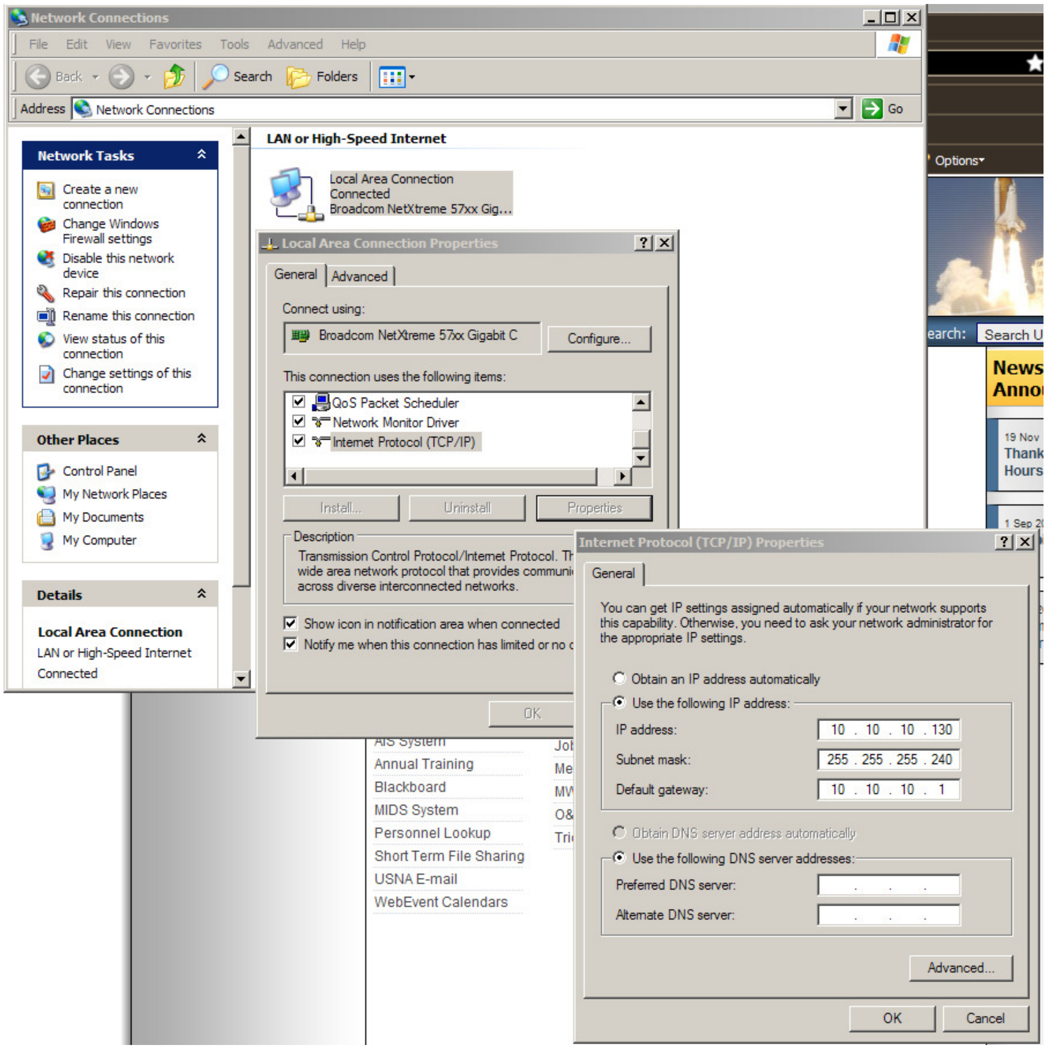

Local Area Connection
Connected
Broadcom NetXtreme 57xx Gig...

Local Area Connection Properties

Connect using:
Broadcom NetXtreme 57xx Gigabit C

This connection uses the following items:
☑ QoS Packet Scheduler
☑ Network Monitor Driver
☑ Internet Protocol (TCP/IP)

Description
Transmission Control Protocol/Internet Protocol. Th
wide area network protocol that provides communi
across diverse interconnected networks.

☑ Show icon in notification area when connected
☑ Notify me when this connection has limited or no c

Internet Protocol (TCP/IP) Properties

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically
● Use the following IP address:
IP address: 10 . 10 . 10 . 130
Subnet mask: 255 . 255 . 255 . 240
Default gateway: 10 . 10 . 10 . 1

○ Obtain DNS server address automatically
● Use the following DNS server addresses:
Preferred DNS server: . . .
Alternate DNS server: . . .

Figure 2- Manually Assigning IP Addresses in Windows

Once the computers are then connected to the switch, the switches enable the communication between the nodes. ICMP echo and echo-replies are an effective way at testing the connections between the two computers, so long as any firewalls are disabled. At this point in the network design, there was no external internet access, making it safe to disable the firewall, (assuming that the two nodes connected to the network were both clean and uninfected).

Virtual LANs

Routers such as the Cisco 3600 series allow for restrictions to be placed upon traffic, dropping and/or allowing traffic to traverse the network. In a network design where several smaller switches are used and connected to one or more routers Virtual LANs are not utilized much. However, when using a Router/Switch such as the Cisco 1811 or other popular models Virtual LANs allow for a port or group of ports to be treated as being an individual switch. Switches operate on the Layer 2 level, and Routers operate on the Layer 3 level, which is to say that switches operate based off MAC addresses where Routers operate on IP addresses. If computers A, B, and C are all connected to the same switch – switch 1 – traffic between the three nodes will never be forwarded to a router. However, if computer C is located on a different switch (whether virtually or physically, in this case switch 2) all traffic from computers A or B destined for computer C will travel through switch 1 to the router, and then be forwarded to switch 2 for delivery to computer C.[4]

This is important, because of the limitations switches encounter operating on the Layer 2 level. Routers are expensive, with the prices increasing with the number of ports on a given router, so it is better to locate (logically, not necessarily physically) similar services on the same LANs/VLANs. A classic example is a MySQL server and Apache Web Server. Placing these two servers in a single LAN/VLAN is a good idea, because these two services will interact together, but will rarely request outside data. Because the type of connections to these services are similar, placing them in their own

[4] Odom, Wendel 2008. CCENT/CCNA ICND1: Official Certification Guide. 696 pages. ISBN:978-1-58720-182-0 Pages 182-190

subnet and LAN/VLAN makes configuring the router's rule-set for traffic filtering easier. Typically an Apache/MySQL server setup will not initialize many, (if any) connections. Also, the database/webserver package will typically only receive requests for connections adhering to a very strict formula. TCP connections on port 80 or 443 are typical for an apache server as these are the HTTP and HTTPS standard port numbers. MySQL operates on port 3306. If these servers are running on a Linux server and/or are not connecting to an ActiveDirectory server, it is likely that these are the only ports that would necessarily require forwarding past the router to the appropriate server. By placing these services in their own isolated subnetwork it is possible to minimized or prevent unauthorized connections to the servers, and also to provide some form of extrusion filtering, by not allowing a compromised node to create future connections in a further attack on the internal network, nor allowing connections to be initialized by the server to an entirely different network in an island-hopping attack.

Configuring Virtual LANs

The Cisco 1811 comes with a Security Device Manager, which is a java-based Graphical User Interface, (GUI), which is a very powerful tool in configuring the Cisco 1811. This tool is especially useful to beginners, as it displays the commands before it issues them to the 1811, allowing the user to eventually learn the format of the commands and later use them on the Command Line Interface, (CLI). Assigning ports to a VLAN was done by clicking checkboxes corresponding to the switch ports and assigning an unique integer value to identify the Virtual LAN. After the settings are applied, the selected port(s) are treated as if they were a physically separate switch connected to the router section of the 1811.

The CLI commands to accomplish this same goal can be accomplished in just a few commands:

- Interface FastEthernet2 //selects FastEthernet2 for port-specific configuration
- Switchport access vlan 102 //assigns FastEthernet2 to VLAN 102.

This assigned the FastEthernet2 port to VLAN 102. These actions may be repeated for any number of FastEthernet ports. The following commands would add FastEthernet3 to VLAN 102. [5]

- Interface FastEthernet3 //selects FastEthernet2 for port-specific configuration
- Switchport access vlan 102 //assigns FastEthernet2 to VLAN 102.

Router Access Control Lists

Router Access Control Lists, (ACLs) provide the ability to apply a filtering rule-set to traffic traveling through a router. The rule-sets are called Access Lists, and there are two types of Access Lists, Standard and Extended. The standard Access Lists can only filter based upon the source IP address,.vice the Extended Access Lists which can filter based upon source and destination IP addresses, port addresses, and even upper-layer protocols such as HTTP, POP3, DNS, etc.[6]

Router Access-Lists can be configured via the GUI on the 1811 router. However, the syntax is fairly simple, and higher-level routers do not all have a GUI, making it worthwhile to learn the CLI method of applying Access-Lists from the beginning. Using the CLI, the administrator may keep a text-file version of the Access List for debugging purposes or for further development. Only one Access List may be assigned to an interface at any time, but one Access List may be assigned to multiple interfaces at any time. Generally, it is wise for the administrator to maintain separate access lists for each interface, as separate access lists are easier to maintain and customize. Following the

[5] Ibid; 10-26

previous examples with the webserver and mysql server, the access list might look something like this.

- Access-list webdb permit tcp any 10.10.10.130 eq 80
- Access-list webdb permit tcp any 10.10.10.130 eq 443
- Access-list webdb permit tcp any 10.10.10.131 eq 3306
- Access-list webdb deny tcp any any
- Access-list webdb deny udp any any [7]

The last two commands technically are unnecessary because there is an implied deny in an access-list at the bottom. For ever packet subject to ACL filtering, the packet is tested by the first rule in the list, and if the first rule does not apply to the specific packet, the next rule is tested, etc. until a rule is found that matches the packet. If none of the rules apply, the deny any any rules take effect, and the packet is dropped. For example a HTTPS session with 10.10.10.130 would not meet the first rule, because the port on HTTPS is 443, and not 80. The packet would then be tested for the second rule, and be permitted, because the rule permits all HTTPS traffic to the web server.

Applying the Access List to an Interface

Writing a secure Access-List is a great step towards securing the network, but many times an inexperienced administrator will forget to bind the access list to an interface. Until the access-list is explicitly applied to an interface, it is a useless file stored on the router.

- Interface vlan102
- ip access-group webdb out

These two commands will bind the access list to the interface so that it is applied to each and every packet will go from the router out to the nodes on VLAN 102, which in this case is the webserver or the database. Currently this only affects traffic directed towards the servers. Any traffic coming

[6] Sedayo, Jeff. 2001. Cisco IOS Access Lists. 272 pages. ISBN: 9781565923850. Pages 231-251.
[7] Odom, Wendel 2008. CCENT/CCNA ICND1: Official Certification Guide. 696 pages. ISBN:978-1-58720-182-0 Pages

from the server is not yet filtered according to any ACL. In order to truly secure the interface VLAN2 the following commands also need to be added.

- Access-list webdb_in permit tcp any 10.10.10.130 0.0.0.0 established
- Access-list webdb_in permit tcp any 10.10.10.130 0.0.0.1 established
- interface vlan102
- ip access-group webdb_in in[8]

Firewall

Firewalls are often regarded as the magic bullet in network security. While this is not entirely true, a powerful firewall is quite possibly the strongest asset a network administrator may have at their disposal. Network firewalls may be software or hardware based, and may be run on an individual leaf-node of a network, or may be placed at the bottleneck of a network to filter and inspect all traffic entering the network from a given point.

There are two main types of firewalls, stateful and stateless, the stateful firewall being more powerful than a stateless firewall. A stateless firewall analyzes the flags on the TCP or UDP packets it receives, and may be configured to block specific ports, IP addresses, or even certain flags in a TCP header, (for example dropping all packets with the SYN flag of a TCP header would disallow all incoming TCP connection requests.) A stateful firewall maintains a table of existing connections and processes packets upon rulesets which can contain conditions to test whether the packets are part of a prior-existing connection, or a new connection. Because stateful firewalls must keep records of existing connections, there are larger processing requirements. This makes stateful firewalls more

231-51

[8] Ibid; 231-51

expensive, but they provide more services. Most stateful firewalls are hardware based.[9]

This research study utilized the PIX 515E firewall, which is a hardware-based firewall and typically configured to run as a stateful firewall. This particular firewall had two Ethernet ports, capable of 100Mbps. Ethernet 0 was designated outside, and Ethernet 1 was designated inside. All traffic departing the network must enter the firewall on Ethernet 1 and depart on Ethernet 0. All traffic entering the network must be passed to Ethernet 0 and then be inspected according the firewall configuration rules.[10] Should an individual packet meet the requirements in the configuration the packet is passed onto the internal network. If not, the packet is "dropped", that is to say that the packet is neither returned nor forwarded to its intended recipient. The firewall, when configured properly can perform many tasks, the following are some of the tasks performed by a PIX 515E:[11]

- Network Address Translation (NAT)
- Port Address Translation (PAT)
- Application filtering
- Protocol Inspection
- Port filtering (source or desination)
- Address Filtering (source or desination)

Network Address Translation (NAT) and Port Address Translation (PAT)

Network Address Translation is the use of one external IP address to match many internal IP addresses. An example of NAT is illustrated below.

[9] Deal, Richard A. 2002. Cisco Pix Firewalls. 537 pages. ISBN:9780072225235. Pages 11-22

[10] Ibid; 12-13

[11] CISCO PIX 515E Security Appliance [Cisco PIX 500 Series Security ApplaInces] – Cisco Systems. 2008. http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5708/ps5709/ps2030/ps4094/product_data_sheet09186a0080091b15.html

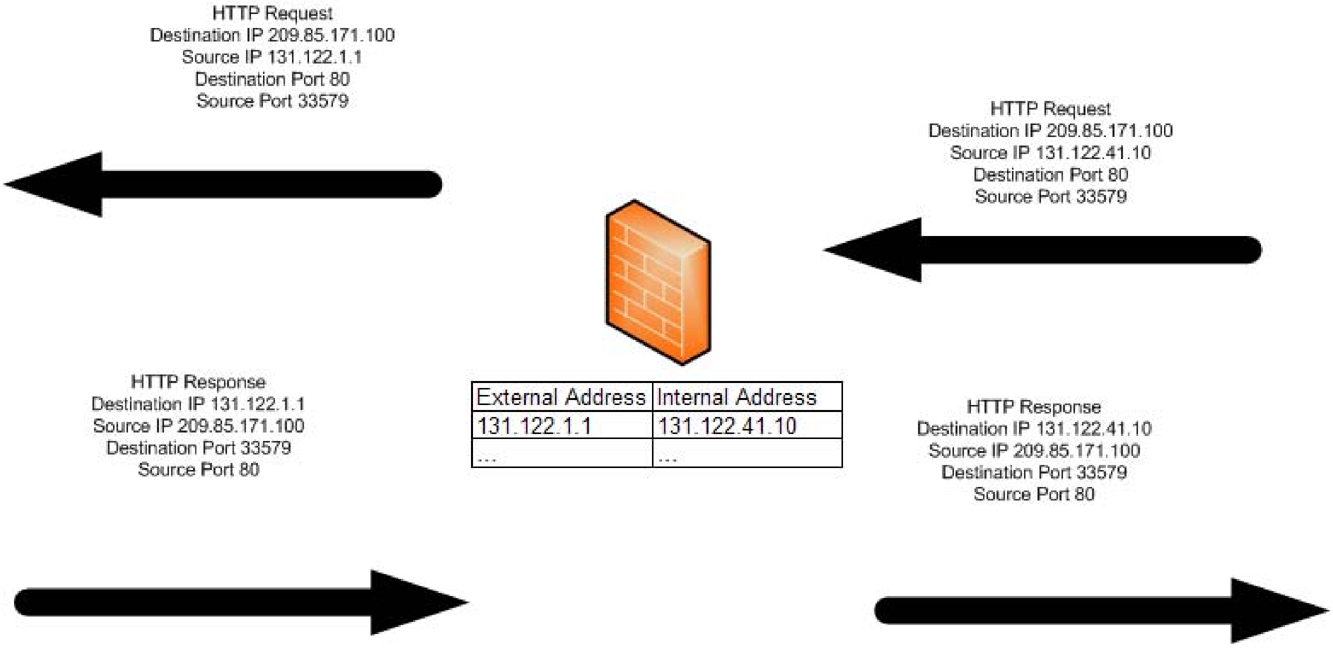

External Address	Internal Address
131.122.1.1	131.122.41.10
...	...

Figure 3- NAT Demonstration

Port Address Translation is similar to NAT, but the port numbers are restricted to smaller pools, each pool belonging to one specific end node on the network. These ports may be statically or dynamically bound to the end nodes.

The PIX 515 is capable of performing NAT and/or PAT, or neither. This means that while the above options may be used, the firewall may be inserted without altering the IP addresses or ports of any packets being altered. In the simulated environment of this research study, the PIX was inserted into the network without any use of NAT or PAT. The routers on both ends of the firewall were configured with static routes to forward all traffic to the firewall, (and beyond).[12]

Initially Configuring the Firewall

There are two ways to configure the PIX firewall. The first is to use the command line interface, (using the console port, and/or the web interface). The other alternative is to use the graphical user interface distributed by Cisco called the PIX Device Manager (PDM). Practically every

[12] Deal, Richard A. 2002. Cisco Pix Firewalls. 537 pages. ISBN: 9780072225235. Pages 116-53

guide online gives instructions for the command line interface, as it appears that few to none of the the system administrators using the PIX utilize the PDM. This project used the command line interface to initially configure the PIX. By connecting the PIX to both the inside and outside routers, and running the setup command, the PIX provides the administrator with an interactive dialog to configure the date and time, inside IP address, IP address of any device running a Trivial File Transfer Protocol Server (TFTP Server), domain name of the network, name of the inside interface, (typically "inside", etc. The outside interface must be configured manually using the following commands:[13]

- interface Ethernet 0   //specify the interface
- name-if outside   //issue a name to the interface ("outside")
- security-level 0   //firewall rules must be applied to any traffic traveling from a
  //lower network security level to a zone with a higher level.
- ip address 192.168.0.1 255.255.255.0   //assign an IP address & subnet mask
- no shutdown   //ensure that the link is not administratively shut down.[14]

The security level is the most important aspect to focus upon in this segment of commands. All packets are inspected and tested whenever traveling from a network with an equal or lower security level than the one the packet is traveling to. On a PIX such as the one in this project which only has two Ethernet ports, this may seem trivial, but on a PIX with three or more Ethernet ports using a Demilitarized Zone (DMZ), this would be more important, because a DMZ contains a network which is trusted less than the inside network, but more than the outside network. This allows for application traffic in the DMZ, (DNS for example) to travel from the DMZ without being inspected, while still allowing for inspection when communications from the DMZ travel to the inside network. By default, the security level of the inside network is configured to be 100, because the security level is a scale from 0 to 100 inclusive and the inside network should be the most trusted of all. The outside interface

[13] How to Configure PIX Firewall Part I. 2008 http://www.secmanager.com/how_to_configure_pix_firewall_part1
[14] PIX Deal, Richard A. 2002. Cisco Pix Firewalls. 537 pages. 9780072225235. Pages 86-103

is set to 0, because the administrator has no control over the rest of the internet, making it the least trusted. [15]

Establishing and Testing Connectivity

Once the firewall has been inserted into the network and assigned IP addresses, it still needs route information before it can actually communicate with any of the other network devices. The simplest way to configure the firewall, at least for testing purposes, is to use static routes. These routes can be specific to an IP address or be a range of IP addresses. [16]

- Route *<interface>* *<ip address>* *<network mask>* *<ip of node to forward to>* *<hops>*
- Route outside 0 0 192.168.0.1 1 //0 0 default for all traffic [17]
- route inside 10.10.10.0 255.255.255.0 10.10.10.1 1 // specifies which IP addresses are
  //on the inside network

At this point, if the PIX and routers have been configured correctly, issuing a ping command from the console port on the PIX firewall to any nodes on the internal network should result in five echo-replies being returned to the router.

Permitting Traffic Flow

In the overall network security implementation, this was the most difficult part of the research. For someone without experience in the network security field, specifically experience with the Cisco IOS for firewalls, this task can prove to be daunting. Even more confusing was that following online tutorials for the basics of firewall configuration I was unable to even get ICMP ping traffic permitted across the network. In a resource constrained environment, the firewall is likely to be the single largest

[15] Ibid; 37-42
[16] Ibid; 101-2
[17] Sedayo, Jeff. 2001. Cisco IOS Access Lists. 272 pages. ISBN: 9781565923850. Page 19

expense. The physical equipment alone can be expensive, plus the time or training required to initially configure the firewall and actually permit desired traffic to traverse the firewall can add up quickly. It might be well worth taking an instructor-led class or some sort of week-long seminar to learn the details of the PIX from the ground up. Noting the many hours spent in the lab that never resulted in success, a real world environment would have lost significant money from network downtime, wages for man-hours, and equipment which might be set aside for a less effective firewall solution after exhausting too much time in futile struggles against one simple error. In the case of this research project, three CCNA certified individuals were unable to determine why the firewall was dropping the traffic.

The Graphical Interface

After running into trouble with the firewall dropping traffic, one CCNA certified individual suggested that I should attempt to utilize the GUI for the PDM, and attempt to configure the firewall using it. The GUI contains a traffic simulator which displays a packet traversing the firewall, and indicates whether the simulated packet would have been permitted or dropped. In the case of the ICMP on this sample network, the simulator confirmed that the packets were being dropped by a firewall rule, and not by other improper configurations.

Lessons Learned

- copy run start – Cisco routers and firewalls have two separate configuration files. The running-configuration is the ruleset which is updated every time a command is issued to the device. When powering up the device, the device loads the startup-config. Without issuing the copy run start command, all changes are lost when the router is reloaded or power is interrupted. Forgetting this command can render hours or days of work useless.

- The hyper-terminal included in Microsoft Windows allows for text to be captured to a text file, or to issue a text file to the device. Issuing the command “terminal length 0” or “pager lines 0” depending upon the device being used before issuing the “show running-config” (device independent) allows the administrator to store, read, edit, and update the configurations using a text editor rather than directly using the command-line.
- Utilize two completely separate networks address spaces for simulations such as this. In order to minimize confusion with rulesets, using the 10 series and the 192.168 series prevents many mistakes.
- Get instructor-led education on the PIX firewall. The firewall is the single most difficult aspect of the network, and without proper configuration, a firewall can render the entire network isolated from the external internet.

Final Accomplishments:

Goals	Methods Used	Notes
Permit communications between two computers	Configure the most basic LAN, (2 nodes connected by a switch)	Accomplished
Connect two separate LANs	Configure a Cisco 3600 Series router to route traffic between the two LANs	Accomplished
Isolate network into compartments to filter extraneous traffic	Establish Virtual LANs.	Accomplished
Only permit standard traffic between VLANs	Configure Access Control Lists to prohibit unexpected traffic traveling between VLANs.	Accomplished
Enable Network Address Translation, protocol inspection, and traffic filtering to shield the internal network from the simulated internet	Introduce a PIX firewall into the network between the 3600 router and the internal LAN, making the connection between the 3600 and the firewall the access-point to the external internet.	Firewall introduced and has connectivity to all nodes on network, however the ruleset is not properly configured, and all traffic is dropped before traversing the firewall.

Table 3 – Final Accomplishments

Advantages of an Independent Study

This independent study provided me with the opportunity to pursue education during my time as a student that would not have otherwise been available to me. The independent study allowed me dedicated access to the equipment I needed to create and maintain the network, and plenty of books and other resources to research the topic at hand. The personnel from the IT Services Department were extremely willing to help in troubleshooting my firewall configurations, and provided valuable insight. The nature of the research project caused me to focus upon an area that most of the Computer Science Department's faculty has limited knowledge in, I was limited only by my ability to scour the internet for information and read various books on the subject. At the completion of this study, I walk away with the knowledge and practical experience required to configure a small-to-medium sized network, and a greater understanding of firewall configuration. While I am by no means an expert in the field of networking after completing this study, I have gained enough appreciation for the networking field to convince me to further my education in network administration and pursue network administration certifications. Although an independent study is not for everyone, I would recommend it to those who are highly self-motivated and capable of dedicating themselves to the long hours of reading, trial and error, and frustration they will experience before the satisfaction of completing the end-result.